

EXACT SIMULATION OF MULTIDIMENSIONAL REFLECTED BROWNIAN MOTION

JOSE BLANCHET* AND
KARTHYEK MURTHY,** *Columbia University*

Abstract

We present the first exact simulation method for multidimensional reflected Brownian motion (RBM). Exact simulation in this setting is challenging because of the presence of correlated local-time-like terms in the definition of RBM. We apply recently developed so-called ε -strong simulation techniques (also known as tolerance-enforced simulation) which allow us to provide a piecewise linear approximation to RBM with ε (deterministic) error in uniform norm. A novel conditional acceptance–rejection step is then used to eliminate the error. In particular, we condition on a suitably designed information structure so that a feasible proposal distribution can be applied.

Keywords: Unbiased sampling; refine until accept/reject; tolerance enforced simulation; acceptance–rejection sampling; ε -strong simulation; iterative algorithm; intersection layer

2010 Mathematics Subject Classification: Primary 65C05
Secondary 60J60; 60J65

1. Introduction

This paper is a contribution to the theory of exact simulation for stochastic differential equations. In particular, we present the first exact simulation algorithm for multidimensional reflected Brownian motion (RBM).

Multidimensional RBM was introduced by Harrison and Reiman [10] and it figures prominently in stochastic operations research. It turns out that RBM approximates the workload at each station in so-called generalized Jackson networks, which are comprised of single-server queues connected via Markovian routing. The approximation holds in heavy traffic (that is, as the system approaches 100% utilization) and it is applicable in great generality (assuming only a functional central limit theorem for the arrival process and the service requirements at each station; see, for example, [6] and [14]). Following [10], we refer to a d -dimensional stochastic process $(Y(t) : t \geq 0)$ that satisfies the following properties as an RBM:

- $Y(\cdot)$ is a Markov process with stationary transition probabilities, continuous sample paths taking values in the nonnegative orthant of \mathbb{R}^d ;
- $Y(\cdot)$ behaves in the interior of a positive orthant like a d -dimensional Brownian motion (either standard or with a constant drift and diffusion matrix);
- $Y(\cdot)$ reflects instantaneously at the boundary of the positive orthant;

Received 12 June 2017.

* Current address: Management Science and Engineering, Stanford University, 475 Via Ortega, Stanford, CA 94305, USA. Email address: jose.blanchet@stanford.edu

** Postal address: Department of Industrial Engineering & Operations Research, Columbia University, S. W. Mudd Building, 500 W. 120 Street, New York, NY 10027, USA.

- the direction of reflection anywhere on the boundary surface where the i th component $Y_i = 0$ is the i th column of the $d \times d$ reflection matrix \mathbf{R} . It is required that \mathbf{R} is of the form $\mathbf{R} = \mathbf{I} - \mathbf{Q}^\top$, where \mathbf{Q} is a nonnegative $d \times d$ matrix with 0s on the diagonal and spectral radius strictly smaller than unity.

RBM as a solution of the Skorokhod problem. Let $(\mathbf{X}(t) : t \geq 0)$ denote a d -dimensional Brownian motion that behaves in a similar manner to RBM $\mathbf{Y}(\cdot)$ in the interior of the positive orthant. Then it is well known that the RBM $\mathbf{Y}(\cdot)$, defined above, can be represented as

$$\mathbf{Y}(t) = \mathbf{X}(t) + \mathbf{R}\mathbf{L}(t) \tag{1.1}$$

with $\mathbf{Y}(t) = (Y_1(t), \dots, Y_d(t))^\top$ and $\mathbf{L}(t) = (L_1(t), \dots, L_d(t))^\top$ satisfying

- (i) $Y_i(t) \geq 0$ for all $t \geq 0$;
- (ii) $L_i(t)$ is nondecreasing in t , and $L_i(0) = 0$;
- (iii) $\int_0^t Y_i(s) dL_i(s) = 0$

for each $t \geq 0$ and $i = 1, \dots, d$ (see, for example, [6] and [10]). We call $\mathbf{X}(\cdot)$ the driving (or free) process and $\mathbf{Y}(\cdot)$ the reflected process. The map S that takes $\mathbf{X}(\cdot)$ to $\mathbf{Y}(\cdot)$ in (1.1) is referred to as the *Skorokhod map*. Item (iii) above simply states that the process $L_i(t)$ increases only at those times t , where $Y_i(t) = 0$. Due to this property, the process $L_i(\cdot)$ behaves like the local time of Brownian motion at the origin. Consequently, the term $\mathbf{R}\mathbf{L}(t)$ appearing in (1.1) is not a standard ‘drift’ term, and cannot be dealt with using change of measure techniques as in [2], [3], or [12].

All the generic exact simulation techniques for diffusions are based on acceptance–rejection, after applying Girsanov’s transformation. The difficulty in applying acceptance–rejection in the multidimensional RBM setting is that there is no natural proposal distribution that can be used to ‘dominate’ the target process directly. In particular, multidimensional RBM is not absolutely continuous with respect to any natural process that is easily simulatable. Note that in one dimension one can simulate RBM directly by keeping track of the running maximum of the driving Brownian motion, and so these challenging issues arise only in dimensions greater than one. Simulation techniques for one-dimensional reflected processes have been studied in [9].

Our contributions. This paper is dedicated to the proof of the following result. Let $\mathbf{Y}(\cdot)$ denote the multidimensional RBM in (1.1).

Theorem 1.1. *Given a deterministic time $T \in (0, 1)$, it is possible to simulate $\mathbf{Y}(T)$ without any bias.*

An obstacle to naively using the traditional acceptance–rejection algorithm (see, for example, [1]) in the simulation of diffusions is that the probability density from which we want to sample is typically unknown. In our setting, while the probability density of $\mathbf{Y}(T)$ itself may be unknown, we propose to simulate enough information about the RBM $\mathbf{Y}(\cdot)$, and perform acceptance–rejection sampling for the probability density of $\mathbf{Y}(T)$ conditional on the simulated filtration. Ideally, the simulated information set should be a collection of random variables such that the probability density of $\mathbf{Y}(T)$ conditional on the simulated information, denoted here by f , is computable. If obtaining such a computable conditional density f is feasible, then one can easily perform an acceptance–rejection step of the form

$$V < \frac{f(\mathbf{Z})}{Cg(\mathbf{Z})} =: L(\mathbf{Z}), \tag{1.2}$$

where g is a suitable proposal density from which proposal samples \mathbf{Z} are drawn, V is an independently generated random variable distributed uniformly in $[0, 1]$, and C is a suitable scaling constant. While this line of thought is interesting, a key difficulty arises from the fact that there is no easily simulatable information structure such that the density of $\mathbf{Y}(T)$ conditional on the simulated information is exactly computable.

To overcome this difficulty, we introduce a novel sampling scheme that we call *refine until accept/reject*, which relaxes the requirement that the conditional density f is known exactly. The key observation behind this algorithm is that in order to accept the proposal \mathbf{Z} , we simply need to decide if (1.2) holds; *we do not need to know the right-hand side of (1.2) exactly*. So, instead of having direct access to the probability density of \mathbf{Y} conditional on \mathcal{I} , if we can simply obtain an approximation to the right-hand side of (1.2) that ensures (1.2) holds, we can accept the proposed sample without incurring any sampling error. We present this idea clearly in a stylized setting in Section 2.1 along with an outline of its applicability to the simulation of RBM in Section 2.2. We use Section 3 to fully present our algorithm for exactly simulating multidimensional RBM. Our algorithm relies on the recently developed ε -strong simulation (also known as tolerance-enforced simulation) techniques of [4] to first derive an approximation of the RBM, which is then used to make one of the following decisions: accept, reject, (or) refine the approximation of the right-hand side in (1.2) until either the proposal can be conclusively accepted or rejected.

We wish to conclude this introduction with a critical discussion of our main result. We do believe that the conditional acceptance–rejection strategy introduced here is of significant value as it addresses an important open problem (exact sampling of multidimensional RBM). Nevertheless, we must recognize that the algorithm, in its current form, is mostly of theoretical interest. Unfortunately, in Section 4 we identify that the expected running time of the algorithm is infinite. While we are investigating strategies to mitigate this problem, we feel that the nucleus of our sampling algorithm, namely *refine-until-accept/reject*, might propel further research in exact sampling of various stochastic processes in addition to the search for efficient sampling algorithms for simulating multidimensional RBM.

2. Overview of the sampling scheme

We first introduce some notational conventions. Throughout the paper we consider the driving (free) stochastic process $\mathbf{X}(\cdot)$ to be a standard Brownian motion in d -dimensions, which we write as $\mathbf{X}(\cdot) = \mathbf{B}(\cdot)$. The reflected process $\mathbf{Y}(\cdot)$ in (1.1) is referred to as the RBM. The extension of our development to the case in which $\mathbf{X}(\cdot)$ is a Brownian motion with constant drift and diffusion coefficients is straightforward. As mentioned in the introduction, the map S that takes $\mathbf{X}(\cdot)$ to $\mathbf{Y}(\cdot)$ in (1.1) is referred to as the Skorokhod reflection map.

While all the variables and stochastic processes taking values in \mathbb{R}^d for $d > 1$ are in bold italic, their one-dimensional counterparts are not. For example, if $\mathbf{B}(\cdot)$ denotes the Brownian motion in multiple dimensions, then $B(\cdot)$ is to be understood as one-dimensional Brownian motion.

2.1. Refine until accept/reject sampling scheme: an introduction

In this section we restrict our attention to the following sampling problem to which our RBM simulation problem is later shown to be reduced. Let Δ and Y be two independent random variables, and $W = Y + \Delta$. For simplicity, we assume that the probability density of Δ , denoted by $f_\Delta(\cdot)$, is continuous on its support which is given by the interval $[-a, a]$ for some $a > 0$. Consecutively, $\sup_{x \in [-a, a]} f_\Delta(x) \leq C$ for some $C \in (0, \infty)$. We assume that Y is also

supported on $[-a, a]$ with an arbitrary distribution from which we do not know how to sample exactly. Our objective in this section is to obtain samples from the distribution of $W = Y + \Delta$.

Given $Y = y$, the density of W evaluated at w is simply $f_\Delta(w - y)$. Moreover, the support of such a density is contained in the interval $[-2a, 2a]$. So, in order to simulate W conditional on y , we could propose W uniformly distributed in $[-2a, 2a]$ and accept if

$$V < C^{-1} f_\Delta(W - y), \tag{2.1}$$

where $V \sim \text{uniform}(0, 1)$ (uniform in $(0, 1)$), and independent of W and $Y = y$.

The key observation is that in order to accept W we simply need to decide if (2.1) holds; *we do not actually need to know the value of y* . So, instead of having direct access to Y , there are settings, as we demonstrate in our RBM simulation, where we might know $\{Y^{\varepsilon_n}\}_{n \geq 1}$, independent of Δ , which converges to Y ; say $\|Y^{\varepsilon_n} - Y\| \leq \varepsilon_n$ for some $\varepsilon_n \rightarrow 0$ as $n \rightarrow \infty$. Then under modest continuity properties of $f_\Delta(\cdot)$, for instance, say $|f_\Delta(x) - f_\Delta(x')| \leq K \|x - x'\|$, we can accept W if

$$V \leq C^{-1} f_\Delta(W - Y^{\varepsilon_n}) - KC^{-1}\varepsilon_n, \tag{2.2}$$

or reject W if

$$V \geq C^{-1} f_\Delta(W - Y^{\varepsilon_n}) + KC^{-1}\varepsilon_n. \tag{2.3}$$

Since $\varepsilon_n \rightarrow 0$ and $V = C^{-1} f_\Delta(W - Y)$ has zero probability of occurring, one must be able to eventually decide whether to accept or reject. As the outlined sampling procedure does one of the following – accept/reject the proposed sample of W , or seek for a refinement of Y^{ε_n} until the proposal can be accepted or rejected – we call the sampling procedure *refine until accept/reject*.

It is useful to remember the following necessary requirements for the procedure underlying (2.1)–(2.3) to yield exact samples of $W = Y + \Delta$.

- (R1) The probability density of Δ , denoted by $f_\Delta(\cdot)$, has bounded support, and is Lipschitz continuous; that is, there exists $K > 0$ such that $|f_\Delta(x) - f_\Delta(x')| \leq K \|x - x'\|$ for all x, x' .
- (R2) Conditional on $Y = y$ and all the information simulated to obtain Y^{ε_n} , the probability density of W evaluated at w is simply $f_\Delta(w - y)$.

2.2. An outline of the application of refine-until-accept/reject for multidimensional RBM

Revisiting our objective of exact sampling of RBM, our plan is to apply the sampling strategy in Section 2.1 by introducing a suitable conditioning. For this purpose, we use the following key facts about multidimensional RBM. First, the fact that if the driving process is Brownian motion then, for fixed T , $\mathbb{P}(Y_i(T) = 0) = 0$ for any $i \in \{1, \dots, d\}$. In addition, since $Y(\cdot)$ is continuous, there exists a $\delta > 0$ and an interval $(T_{\text{left}}, T_{\text{right}}]$ which contains T , satisfying $Y_i(s) > \delta$ for all $i \in \{1, \dots, d\}$, and, therefore,

$$Y(s) = Y(T_{\text{left}}) + X(s) - X(T_{\text{left}}) \quad \text{for all } s \in (T_{\text{left}}, T_{\text{right}}].$$

In other words, the interval $(T_{\text{left}}, T_{\text{right}}]$ is such that the RBM $Y(t)$ does not hit the reflecting boundary anywhere during $t \in (T_{\text{left}}, T_{\text{right}}]$ and, consecutively, $L(s) - L(T_{\text{left}}) = 0$ for all $s \in (T_{\text{left}}, T_{\text{right}}]$.

So, our plan is to first simulate enough information about $X(\cdot)$ (that is, the driving Brownian motion) so that, conditional on such information, we have the representation

$$Y(T) = Y(T_{\text{left}}) + (X(T) - X(T_{\text{left}})) =: Y(T_{\text{left}}) + \Delta \tag{2.4}$$

for a suitable T_{left} identified from the simulated information. Naturally, we identify $Y(T)$ and $Y(T_{\text{left}})$, respectively, with the variables W and Y introduced in the abstract setting discussed previously in Section 2.1. Therefore, our objective is to simulate just enough information so that, conditioned on the simulated information,

(R1') the probability density of the Brownian increment $\Delta := X(T) - X(T_{\text{left}})$, denoted by $f_{\Delta}(\cdot)$, has bounded support, and is Lipschitz continuous, and

(R2') the probability density of $Y(T)$, evaluated at w , is simply $f_{\Delta}(w - Y(T_{\text{left}}))$.

The requirements (R1') and (R2') mirror the earlier requirements (R1) and (R2) in the abstract setting in Section 2.2. Once these requirements are met, we can follow the logic in Section 2.1 to develop a refine-until-accept/reject sampler for obtaining samples from the distribution of $Y(T)$. Thus, the proposed algorithm for exact sampling of $Y(T)$ can be roughly divided into two steps:

- a preconditioning step where we simulate enough information to arrive at representation (2.4), and
- exploiting representation (2.4) arrived at in the preconditioning step, we perform refine-until-accept/reject to obtain samples of $Y(T)$.

2.2.1. Overview of the preconditioning step. In order to sample enough information which will enable us to obtain representation (2.4) along with satisfying the above two requirements, we use another important property of the Skorokhod map S , namely, S is Lipschitz continuous as a function of the driving process in the uniform norm over the time interval $[0, 1]$. Consequently, to identify T_{left} we use so-called ε -strong simulation techniques, also known as tolerance-enforced simulation, which allows us to simulate $X^{\varepsilon}(\cdot)$ piecewise linear and guaranteed to be within ε -close in uniform norm to $X(\cdot)$. This construction is, conceptually at least, not complicated. There are several methods that can be applied to the same end: based on wavelets [5], localization using stopping times [7], or tracking jointly the maxima and minima on dyadic intervals [4]. We have chosen to use the latter construction, thereby ultimately obtaining $(T_{\text{left}}, T_{\text{right}})$ as a dyadic interval (that is, $T_{\text{left}} = i2^{-N}$ and $T_{\text{right}} = j2^{-N}$ for some $0 \leq i < j \leq 2^N$ and $N > 0$). The reason for choosing the construction in [4] is because it allows us to recursively develop more refined approximations $X^{\varepsilon'}$ for desired $\varepsilon' < \varepsilon$ while preserving the conditional independence of Δ and $Y(T_{\text{left}})$ given all the information required to conclude that $L(T_{\text{right}}) - L(T_{\text{left}}) = 0$. We refer the reader to Section 3.1.1 for an overview of the algorithm of [4] that allows us to obtain the desired piecewise linear approximation $X^{\varepsilon}(\cdot)$ for the driving Brownian motion.

The Skorokhod problem is easy to solve for piecewise linear input X^{ε} , because in such a case the solution to the Skorokhod problem, denoted by $(Y^{\varepsilon}, L^{\varepsilon})$, is piecewise linear as well, and the gradients can be obtained by solving linear systems based on (1.1) (see Section 3.1.2 for an explicit algorithm). Since the piecewise linear approximation $Y^{\varepsilon_n}(\cdot)$ to RBM $Y(\cdot)$ can be identified explicitly for a computable $\{\varepsilon_n\}_{n \geq 1}$ such that $\varepsilon_n \rightarrow 0$ as $n \rightarrow \infty$, the Lipschitz continuity of $Y = S(X)$ as a function of X , combined with the approximation X^{ε_n} , and the fact that $Y(T)$ must be strictly positive coordinatewise, can eventually be used to identify T_{left} used in the additive representation (2.4). See Section 3.1 for details.

2.2.2. Requirements for the refine-until-accept/reject step. Once we arrive at representation (2.4), we can use the refine-until-accept/reject algorithm introduced in Section 2.1 to obtain samples of $Y(T)$. In order to be able to do this, we need to guarantee that the requirements

(R1') and (R2') mentioned earlier are met. Our construction of $X^\epsilon(\cdot)$, as indicated earlier, based on [4] will give rise to a conditional density for Δ , denoted by $f_\Delta(\cdot)$, which is expressed as an infinite series. So, the Lipschitz continuity of $f_\Delta(\cdot)$ used in (2.2) and (2.3) is obtained by means of some careful estimates. Consequently, as we shall see in Section 3.2, we will be able to implement the basic refine-until-accept/reject strategy underlying (2.1)–(2.3).

We now provide more specific details behind our sampling methodology and point to future relevant sections where details are fully fleshed out.

3. The exact sampling scheme

We first describe essential components of the preconditioning step (such as ϵ -strong approximation techniques) before presenting the refine-until-accept/reject sampler for $Y(T)$.

3.1. The preconditioning step

As mentioned in Section 2.2, the objective of the preconditioning step is to simulate just enough information in order to obtain representation (2.4) while satisfying requirements (R1') and (R2').

3.1.1. *Generating an ϵ -strong approximation for the driving Brownian motion.* Here, we first provide a brief description of the ϵ -strong algorithm of [4] that simulates a piecewise linear approximation to the one-dimensional standard Brownian motion $B(\cdot)$. The algorithm iteratively generates a sequence of pairs of piecewise constant dominating processes, $\{B_n^\uparrow(t) : t \in [0, 1]\}$ and $\{B_n^\downarrow(t) : t \in [0, 1]\}$, that satisfy the following properties. For all $t \in [0, 1]$,

$$B_n^\downarrow(t) \leq B_{n+1}^\downarrow(t) \leq B(t) \leq B_{n+1}^\uparrow(t) \leq B_n^\uparrow(t),$$

$$\sup_{t \in [0, 1]} |B_n^\uparrow(t) - B_n^\downarrow(t)| \searrow 0 \quad \text{almost surely as } n \nearrow \infty.$$

At every step $n \geq 1$, the algorithm generates information about the Brownian motion $B(\cdot)$ in dyadic intervals $\{(j-1)2^{-n}, j2^{-n}\} : j = 1, \dots, 2^n\}$ conditional on the information available on dyadic intervals from the $(n-1)$ th step. Let $m_{j,n}$ and $M_{j,n}$ denote the extrema of $B(\cdot)$:

$$m_{j,n} := \inf\{B(t) : t \in ((j-1)2^{-n}, j2^{-n})\}, \quad M_{j,n} = \sup\{B(t) : t \in ((j-1)2^{-n}, j2^{-n})\}.$$

During the n th iteration, the ϵ -strong algorithm simulates the following random quantities for each dyadic interval (indexed by $j = 1, \dots, 2^n$):

- an interval that contains the minimum: $L_{j,n}^\downarrow$ and $L_{j,n}^\uparrow$ such that $m_{j,n} \in [L_{j,n}^\downarrow, L_{j,n}^\uparrow]$ and $L_{j,n}^\uparrow - L_{j,n}^\downarrow < 2^{-(n+1)/2}$;
- an interval that contains the maximum: $U_{j,n}^\downarrow$ and $U_{j,n}^\uparrow$ such that $M_{j,n} \in [U_{j,n}^\downarrow, U_{j,n}^\uparrow]$ and $U_{j,n}^\uparrow - U_{j,n}^\downarrow < 2^{-(n+1)/2}$;
- the end-points of Brownian motion: $B((j-1)2^{-n})$ and $B(j2^{-n})$.

Let $\mathcal{I}_{j,n}$ denote the collective information

$$\mathcal{I}_{j,n} := \{L_{j,n}^\downarrow, L_{j,n}^\uparrow, U_{j,n}^\downarrow, U_{j,n}^\uparrow, B((j-1)2^{-n}), B(j2^{-n})\},$$

which is referred to as the *intersection layer* in [4]. Let \mathcal{I} denote the collection of all the intersection layers; at the end of the n th iteration, the collection \mathcal{I} is updated as

$$\mathcal{I} := \{\mathcal{I}_{j,n} : j = 1, \dots, 2^n\}.$$

The $(n+1)$ th step makes use of \mathcal{J} generated in the n th step to arrive at $\{\mathcal{J}_{j,n+1} : j = 1, \dots, 2^{n+1}\}$. Specific details of how these random quantities are simulated can be found in [4]. From the intersection layers $\mathcal{J}_{j,n}$ generated by the algorithm at the n th step, the dominating piecewise constant processes can be formed as

$$B_n^\uparrow(t) = \sum_{j=1}^{2^n} U_{j,n}^\uparrow \mathbf{1}(t \in ((j-1)2^{-n}, j2^{-n})),$$

$$B_n^\downarrow(t) = \sum_{j=1}^{2^n} L_{j,n}^\downarrow \mathbf{1}(t \in ((j-1)2^{-n}, j2^{-n})),$$

where $\mathbf{1}(\cdot)$ is the indicator function. Further, define the following piecewise linear process which shall serve as our approximation for $B(\cdot)$:

$$B_n(t) = \sum_{j=1}^{2^n} [B((j-1)2^{-n}) + 2^n(B(j2^{-n}) - B((j-1)2^{-n}))(t - (j-1)2^{-n})], \quad (3.1)$$

which is just a linear interpolation of the points $\{B((j-1)2^{-n}) : j = 1, \dots, 2^n\}$ over the dyadic intervals in $[0, 1]$. Note that all the random variables used in the construction of $B_n^\uparrow(\cdot)$, $B_n^\downarrow(\cdot)$, and $B_n(\cdot)$ are available in \mathcal{J} , and can be simulated on a personal computer without any discretisation error. It was proved in [4] that the dominating processes $B_n^\uparrow(\cdot)$ and $B_n^\downarrow(\cdot)$ have the following convergence behavior:

$$\varepsilon_n := \sup_{t \in [0,1]} |B_n^\uparrow(t) - B_n^\downarrow(t)| = \max_{1 \leq j \leq 2^n} |U_{j,n}^\uparrow - L_{j,n}^\downarrow| \searrow 0,$$

$$\mathbb{E} \left[\int_0^1 |B_n^\uparrow(t) - B_n^\downarrow(t)| dt \right] = O(2^{-n/2}) \quad \text{as } n \rightarrow \infty.$$

To generate a piecewise linear approximation of the d -dimensional Brownian motion $B(\cdot) = (B_1(\cdot), \dots, B_d(\cdot))$, we generate approximating processes $B_{n,i}(\cdot)$, $B_{n,i}^\uparrow(\cdot)$, and $B_{n,i}^\downarrow(\cdot)$ independently for each one-dimensional Brownian motion $B_i(\cdot)$ as explained above, and use

$$B_n(t) = (B_{n,1}(t), \dots, B_{n,d}(t)), \quad t \in [0, 1],$$

as a piecewise linear approximation for $B(\cdot)$. Similar to the one-dimensional case, the simulated information is stored in the intersection layers $\mathcal{J} = \{\mathcal{J}_{j,n}^i : j = 1, \dots, 2^n, i = 1, \dots, d\}$; here, $\mathcal{J}_{j,n}^i$ simply denotes the intersection layer simulated to generate the approximation to the i th component $B_i(t)$ of the driving Brownian motion $B(\cdot)$. As in the one-dimensional case, we use ε_n to denote the error in approximation at the n th step:

$$\varepsilon_n := \sup\{|B_{n,i}^\uparrow(t) - B_{n,i}^\downarrow(t)| : t \in [0, 1], i = 1, \dots, d\}. \quad (3.2)$$

3.1.2. Generating an ε -strong approximation of RBM. Given a linear path $(\mathbf{x}(t) : t \in [t_0, t_1])$ specified by initial conditions $\mathbf{x}(t_0) \in \mathbb{R}^d$ and $\dot{\mathbf{x}}(t) = \mathbf{m} \in \mathbb{R}^d$ for $t \in (t_0, t_1)$, we explain in this section how to identify the reflected path $\mathbf{y}(\cdot) = S(\mathbf{x}(\cdot))$ that solves the Skorokhod problem in (1.1) with $(\mathbf{x}(t) : t \in [t_0, t_1])$ as the path of the driving (free) process. Once we know how to solve (1.1) for a linear path in the interval (t_0, t_1) , it is straightforward to iteratively solve for any piecewise linear path of the driving (free) process.

As the slope of each component of the driving path $\mathbf{x}(\cdot)$ is fixed in the interval $[t_0, t_1]$, the slope of \mathbf{y} at time $t \in (t_0, t_1)$, denoted by $\dot{\mathbf{y}}(t)$, is obtained from

$$\dot{\mathbf{y}}(t) = \mathbf{m} + \mathbf{R}\mathbf{z},$$

where, as per conditions (i)-(iii) in Skorokhod problem (1.1), $\mathbf{y}(t)$ and $\mathbf{z} \in \mathbf{R}^d$ should satisfy

$$\mathbf{y}(t) \geq \mathbf{0}, \quad \mathbf{z} \geq \mathbf{0}, \quad \mathbf{y}(t)\mathbf{z} = \mathbf{0}.$$

As a component z_i of $\mathbf{z} = (z_1, \dots, z_d)$ is nonzero only when the respective component $y_i(t)$ of $\mathbf{y}(t)$ is 0, it is useful to dynamically keep track of which of the components of $\mathbf{y}(t) = (y_1(t), \dots, y_d(t))$ are 0. We accomplish this algorithmically by letting $C := \{i : y_i(t) = 0\}$. In addition, we use \mathbf{R}_C to denote the submatrix of \mathbf{R} formed by letting $\mathbf{R}_C = [R_{i,j}]_{i,j \in C}$; similarly, let $\mathbf{z}_C = [z_i]_{i \in C}$, $\dot{\mathbf{y}}_C(t) = [\dot{y}_i(t)]_{i \in C}$ denote the vectors formed by entries restricted to indices from the set C . As \mathbf{z}_C is the minimal nonnegative vector that maintains $\mathbf{y}_C(t) = 0$ (see [11]), it is obtained by solving a linear program as in Algorithm 1 when any of the components of \mathbf{y} hit 0.

Algorithm 1. An algorithm to solve the Skorokhod problem (1.1) for a linear path $(\mathbf{x}(t) : t \in [t_0, t_1])$ with componentwise constant slopes \mathbf{m} in the interval $[t_0, t_1]$ for given initial condition $\mathbf{y}(t_0)$.

procedure approxRBM($t_0, t_1, \mathbf{y}(t_0), \mathbf{m}$)

Initialize the time data structure $\vec{t} = [t_0]$ and path space data structure $\vec{\mathbf{y}} = [\mathbf{y}(t_0)]$

Initialize $\mathbf{y} = \mathbf{y}(t_0)$, $C = \{i : y_i = 0\}$, $t' = t_0$, $\mathbf{z} = \mathbf{0}$.

while $t' < t_1$ **do**

Solve the linear program $\min_{\mathbf{z}_C, \dot{\mathbf{y}}_C} \mathbf{1}^\top \mathbf{z}$,

subject to $\mathbf{z}_C \geq \mathbf{0}$, $\dot{\mathbf{y}}_C = \mathbf{m}_C + \mathbf{R}_C \mathbf{z}_C$, $\dot{\mathbf{y}}_C \geq \mathbf{0}$.

Update $\dot{\mathbf{y}} = \mathbf{m} + \mathbf{R}\mathbf{z}$

$$t' = \min \left\{ t_1, \min_{i: \dot{y}_i < 0} \left(t_0 - \frac{y_i}{\dot{y}_i} \right) \right\},$$

$$\mathbf{y} = \mathbf{y} + \dot{\mathbf{y}}(t' - t_0) \text{ and}$$

$$C = \{i : y_i = 0\}.$$

Append the time and path space data structures with the latest entries as in

$$\vec{t} \leftarrow [\vec{t} \quad t'] \quad \text{and} \quad \vec{\mathbf{y}} \leftarrow [\vec{\mathbf{y}} \quad \mathbf{y}].$$

end while

Return $\vec{t}, \vec{\mathbf{y}}$.

Given componentwise constant slopes \mathbf{m} of the driving process $\mathbf{x}(\cdot)$ in the interval (t_0, t_1) and the initial condition $\mathbf{y}(t_0) = \mathbf{y}$, Algorithm 1 returns a vector \vec{t} and a matrix $\vec{\mathbf{y}}$ of dimension $d \times \text{size}(\vec{t})$, where $\text{size}(\vec{t})$ is the dimension of \vec{t} . The data structures \vec{t} and $\vec{\mathbf{y}}$ returned by Algorithm 1 can be used to construct the piecewise linear path $(\mathbf{y}(t) : t \in [t_0, t_1])$ as follows. If the i th entry of \vec{t} is t_i and the i th column of the matrix $\vec{\mathbf{y}}$ is \mathbf{y}_i , then linear interpolation of the points (t_i, \mathbf{y}_i) yields the reflected path $(\mathbf{y}(t) : t \in [t_0, t_1])$.

3.1.3. *Detecting the interval* ($T_{\text{left}}, T_{\text{right}}$]. Recall that the objective of the preconditioning step is to simulate just enough information about the RBM so that we arrive at representation (2.4). As explained in Section 2.2, we propose to achieve this by detecting an interval ($T_{\text{left}}, T_{\text{right}}$] containing T such that the RBM stays in the interior of positive orthant without hitting the reflecting boundary anywhere in that interval. In order to accomplish this, we first make the following observations.

- (A) It is well known that the Skorokhod map S is Lipschitz continuous (with respect to the uniform metric on the path space $C([0, 1]; \mathbb{R}^d)$) with Lipschitz constant $K_1 := (1 - \alpha)^{-1}$; here, $\alpha \in (0, 1)$ denotes the spectral radius of the matrix $\mathbf{Q} = \mathbf{I} - \mathbf{R}$. (see [10]). Therefore, if we solve (1.1) with $X(\cdot) = \mathbf{B}_n(\cdot)$ as the driving (free) process, the corresponding reflected process $\mathbf{Y}_n(\cdot) := S(\mathbf{B}_n)(\cdot)$ satisfies

$$\|\mathbf{Y}_n - \mathbf{Y}\| := \sup_{t \in [0, 1], i = 1, \dots, d} |Y_{n,i}(t) - Y_i(t)| < K_1 \varepsilon_n, \tag{3.3}$$

where $\varepsilon_n := \sup\{B_{n,i}^\uparrow(t) - B_{n,i}^\downarrow(t) : t \in [0, 1], i = 1, \dots, d\}$, and $B_{n,i}^\uparrow(\cdot), B_{n,i}^\downarrow(\cdot)$ are, respectively, the simulated upper and lower bounding processes of the i th component of the driving Brownian motion $\mathbf{B}(\cdot)$. Since

$$\|\mathbf{B}_n(t) - \mathbf{B}(t)\| = \sup\{|B_{n,i}(t) - B_i(t)| : t \in [0, 1], i = 1, \dots, d\} \leq \varepsilon_n,$$

(3.3) follows as a simple consequence of the Lipschitz continuity of the Skorokhod map.

- (B) We say that $a := \min_{i=1, \dots, d} Y_i(T) > 0$; that is, the RBM $\mathbf{Y}(\cdot)$ lies in the interior of positive orthant at time T . Then, due to the continuity of Brownian motion paths, there exists a random interval $(T - t', T + t')$ where $|B_i(t) - B_i(T)| < a$ for all $t \in (T - t', T + t')$, and, consequently,

$$\mathbf{Y}(t) - \mathbf{Y}(T) = \mathbf{B}(t) - \mathbf{B}(T), \quad \mathbf{Y}(t) = \mathbf{Y}(T) + (\mathbf{B}(t) - \mathbf{B}(T)) > a\mathbf{1} - a\mathbf{1} = \mathbf{0}$$

for all $t \in (T - t', T + t')$. In other words, the RBM $\mathbf{Y}(t)$ lies in the interior of positive orthant for every $t \in (T - t', T + t')$ where $|B_i(t) - B_i(T)| < a$. In order to make use of this observation, recall that our constructed piecewise constant upper and lower bounding processes $B_{n,i}^\uparrow$ and $B_{n,i}^\downarrow$ satisfy $B_{n,i}^\uparrow(\cdot) - B_{n,i}^\downarrow(\cdot) < \varepsilon_n$ for each component i , and, consequently,

$$|B_i(t) - B_i(s)| < \varepsilon_n, \quad t, s \in [(j - 1)2^{-n}, j2^{-n}]$$

for all $i = 1, \dots, d$ and $j = 1, \dots, 2^n$. As a result, if $\min_{i=1, \dots, d} Y_i(T) > \varepsilon_n$ then as $|B_i(t) - B_i(T)| < \varepsilon_n$ for every t in the dyadic interval containing T , the RBM $\mathbf{Y}(\cdot)$ does not hit the reflecting boundary anywhere in the specific dyadic interval $[(j - 1)2^{-n}, j2^{-n}]$ containing T .

If we know that ε_n is small enough such that every component of the RBM satisfies $Y_i(T) \geq \varepsilon_n$, then from observation (B) above, the RBM $\mathbf{Y}(t)$ stays in the interior of positive orthant for every t in the dyadic interval $((j - 1)2^{-n}, j2^{-n}]$ containing T ; consecutively, we can declare the corresponding interval $((j - 1)2^{-n}, j2^{-n}]$ as $(T_{\text{left}}, T_{\text{right}}]$. However, since we do not know $\mathbf{Y}(T)$, the immediate objective is to figure out how to guarantee that $Y_i(T)$ is indeed larger than ε_n for every $i = 1, \dots, d$. From the Lipschitz continuity in observation (A), if ε_n is small enough so that $Y_{n,i}(T) = S(\mathbf{B}_n)(T) > (K_1 + 1)\varepsilon_n$ for some n , then $Y_i(T) > \varepsilon_n$.

Since $Y(t)$ lies in the interior of the positive orthant almost everywhere, we will indeed have $Y_{n,i}(T) > (K_1 + 1)\varepsilon_n$ for suitably small approximation error ε_n . Now define

$$N := \inf\{n \geq 1 : Y_{n,i}(T) > (K_1 + 1)\varepsilon_n, i = 1, \dots, d\} \quad \text{and} \quad \delta := \varepsilon_N.$$

Recall that $\varepsilon_n := \sup\{|B_{n,i}^\uparrow(t) - B_{n,i}^\downarrow(t)| : t \in [0, 1], i = 1, \dots, d\}$. The preconditioning procedure for detecting the interval $(T_{\text{left}}, T_{\text{right}}]$ that simulates approximations to Brownian motion and RBM until the stopping time N is summarized in Algorithm 2.

Algorithm 2. Preconditioning step to arrive at representation (2.4) (achieved mainly via ε -strong approximation of the underlying stochastic processes).

procedure preconditioning(T)

Initialize $\mathcal{J} = \emptyset, n = 0, \delta = 1, Y_0(t) = 0, t \in [0, 1]$

while $n = 0$ OR $Y_{n,i}(T) < (K_1 + 1)\delta$ for some $i \in \{1, \dots, d\}$ **do**

Increment $n \leftarrow n + 1$

Simulate the intersection layers $\mathcal{J}_{j,n}^i$ for $j = 1, \dots, 2^n, i = 1, \dots, d$ and the piecewise constant upper and lower bounding processes $B_{n,i}^\uparrow(\cdot), B_{n,i}^\downarrow(\cdot)$, conditional on \mathcal{J}

Form $B_n(\cdot) = (B_{n,1}, \dots, B_{n,d})$ as in (3.1), which serves as a piecewise linear approximation to $B(\cdot)$

Update $\delta = \sup\{|B_{n,i}^\uparrow(t) - B_{n,i}^\downarrow(t)| : t \in [0, 1], i = 1, \dots, d\}$

$$\mathcal{J} = \{\mathcal{J}_{j,n} : j = 1, \dots, 2^n\}$$

Letting $X(\cdot) = B_n(\cdot)$ in (1.1), solve for the reflected process $Y(\cdot)$ using Algorithm 1; call the solution as $Y_n(\cdot)$

end while

Find $J = \{1 \leq j \leq 2^n : (j - 1)2^{-n} < T \leq j2^{-n}\}$

Set $N = n, T_{\text{left}} = (J - 1)2^{-n}$ and $T_{\text{right}} = J2^{-n}$

Return $\mathcal{J}, N, T_{\text{left}}, T_{\text{right}}, \delta$ and $Y^\delta(\cdot) := Y_n(\cdot)$

With this construction, since the RBM $Y(\cdot)$ does not hit the reflecting boundary anywhere in the interval $(T_{\text{left}}, T_{\text{right}}]$, the dynamics of $B(\cdot)$ and $Y(\cdot)$ match those in $(T_{\text{left}}, T_{\text{right}}]$; in particular,

$$Y(t) - Y(T_{\text{left}}) = B(t) - B(T_{\text{left}}) \quad \text{for all } t \in (T_{\text{left}}, T_{\text{right}}],$$

thus resulting in the desired additive representation $Y(T) = Y(T_{\text{left}}) + \Delta$ in (2.4), where the increment $\Delta := B(T) - B(T_{\text{left}})$ is simply the Brownian increment.

3.2. The conditional probability density of Δ for the refine-until-accept/reject sampler

The requirements (R1') and (R2'), listed in Section 2.2 for the implementation of the refine-until-accept/reject exact sampler, necessitate that we know the law of Δ conditional on all the simulated collection of random variables \mathcal{J} . For ease of exposition, we consider the one-dimensional case; the conditional probability density of the Brownian increment Δ in d -dimensions, denoted by $f_\Delta(\cdot)$, is given simply by the product form of one-dimensional densities.

At any stage of the algorithm, all the information simulated about the driving Brownian motion is available in the intersection layers $\mathcal{I} = \{\mathcal{I}_{j,n} : j = 1, \dots, 2^n\}$. From Algorithm 2, recall that J is the index corresponding to the dyadic interval $(T_{\text{left}}, T_{\text{right}}]$ that contains T ; that is, $(J - 1)2^{-N} = T_{\text{left}} < T \leq T_{\text{right}} = J2^{-N}$. For ease of notation, let

$$\begin{aligned} L^\downarrow &:= L_{J,N}^\downarrow - B(T_{\text{left}}), & L^\uparrow &:= L_{J,N}^\uparrow - B(T_{\text{left}}), \\ U^\downarrow &:= U_{J,N}^\downarrow - B(T_{\text{left}}), & U^\uparrow &:= U_{J,N}^\uparrow - B(T_{\text{left}}), \\ l &:= T_{\text{right}} - T_{\text{left}}, & s &:= T - T_{\text{left}}, & v &:= B(T_{\text{right}}) - B(T_{\text{left}}). \end{aligned}$$

Further, let $W(\cdot)$ denote an independent standard Brownian motion on $C[0, 1]$ under measure $\mathbb{P}(\cdot)$. Then, due to the Markov property of $B(\cdot)$, the increment Δ conditional on \mathcal{I} has the following density:

$$f_\Delta(x) dx = \mathbb{P}\left\{W(s) \in dx \mid W(l) = v, \inf_{0 \leq t \leq l} W(t) \in (L^\downarrow, L^\uparrow), \sup_{0 \leq t \leq l} W(t) \in (U^\downarrow, U^\uparrow)\right\}.$$

Note that the support of $f_\Delta(\cdot)$ is $(L^\downarrow, U^\uparrow)$. A closed-form expression for $f_\Delta(\cdot)$ follows from Proposition 5.1 of [4], and is given here:

$$f_\Delta(x) \propto \rho(x) \times \pi(x),$$

where, for any fixed $L^\downarrow, L^\uparrow, U^\downarrow, U^\uparrow, v, s,$ and $l,$

$$\begin{aligned} \pi(x) &:= \exp\left(-\frac{1}{2}\left(x - \frac{s}{l}v\right)^2 \Big/ \frac{s(l-s)}{l}\right), \\ \rho(x) &:= \mathbb{P}\left\{\inf_{0 \leq t \leq l} W(t) \in (L^\downarrow, L^\uparrow), \sup_{0 \leq t \leq l} W(t) \in (U^\downarrow, U^\uparrow) \mid W(s) = x, W(l) = v\right\} \\ &= \gamma_1(x)\gamma_2(x) - \gamma_3(x)\gamma_4(x) - \gamma_5(x)\gamma_6(x) + \gamma_7(x)\gamma_8(x). \end{aligned} \tag{3.4}$$

To define $\gamma_1, \dots, \gamma_8,$ first consider the probability that the Brownian bridge from a to b in the time interval $[0, r]$ stays within (L, U) :

$$\begin{aligned} \gamma(L, U; r, a, b) &:= \mathbb{P}\left\{L < \inf_{0 \leq t \leq l} W(t) \leq \sup_{0 \leq t \leq l} W(t) < U \mid W(0) = a, W(r) = b\right\} \\ &= \left(1 - \sum_{j=1}^\infty (\sigma_j - \tau_j)\right) \mathbf{1}(a, b \in (L, U)), \end{aligned} \tag{3.5}$$

where,

$$\begin{aligned} \sigma_j &:= \exp\left(-\frac{2}{r}((U - L)j + L - a)((U - L)j + L - b)\right) \\ &\quad + \exp\left(-\frac{2}{r}((U - L)j - U + a)((U - L)j - U + b)\right), \\ \tau_j &:= \exp\left(-\frac{2(U - L)j}{r}((U - L)j + a - b)\right) + \exp\left(-\frac{2(U - L)j}{r}((U - L)j + b - a)\right). \end{aligned}$$

Equation (3.5) for $\gamma(L, U; l, a, b)$ is originally from [13]. Now we are ready to define $\gamma_1, \dots, \gamma_8$ mentioned in (3.4):

$$\begin{aligned} \gamma_1(x) &= \gamma(L^\downarrow, U^\uparrow; s, 0, x), & \gamma_2(x) &= \gamma(L^\downarrow, U^\uparrow; l - s, x, v), \\ \gamma_3(x) &= \gamma(L^\uparrow, U^\uparrow; s, 0, x), & \gamma_4(x) &= \gamma(L^\uparrow, U^\uparrow; l - s, x, v), \\ \gamma_5(x) &= \gamma(L^\downarrow, U^\downarrow; s, 0, x), & \gamma_6(x) &= \gamma(L^\downarrow, U^\downarrow; l - s, x, v), \\ \gamma_7(x) &= \gamma(L^\uparrow, U^\downarrow; s, 0, x), & \gamma_8(x) &= \gamma(L^\uparrow, U^\downarrow; l - s, x, v). \end{aligned}$$

To perform acceptance–rejection-type sampling, we need the conditional density f_Δ (of the Brownian increment) to be Lipschitz continuous (see requirement (R1') in Section 2.2). Lemma 3.1 is a step towards establishing this fact.

Lemma 3.1. *There exist positive constants $c_\pi, K_\pi, c_\rho,$ and K_ρ such that, for any fixed $L^\downarrow, L^\uparrow, U^\downarrow, U^\uparrow, v, s$ and $l,$*

$$\begin{aligned} \pi(x) < c_\pi, & \quad |\pi(x) - \pi(y)| < K_\pi|x - y|, \\ \rho(x) < c_\rho, & \quad |\rho(x) - \rho(y)| < K_\rho|x - y|, \end{aligned} \quad \text{for all } x, y \in (L^\downarrow, U^\uparrow).$$

Explicit closed-form expressions for the constants $K_\pi, K_\rho, c_\pi,$ and c_ρ are presented in Appendix A.

Following the representation that $Y(T) = Y(T_{\text{left}}) + \Delta,$ the conditional density of $Y(T),$ given $Y(T_{\text{left}}),$ is given by $f_\Delta(w - Y(T_{\text{left}})),$ which is supported on $(Y(T_{\text{left}}) + L^\downarrow, Y(T_{\text{left}}) + U^\uparrow) \subset (Y(T_{\text{left}}) - \delta, Y(T_{\text{left}}) + \delta).$ As the unknown $Y(T_{\text{left}})$ differs from $Y^\delta(T_{\text{left}})$ only by $\pm K_1\delta,$ the support of $Y(T)$ conditional on the simulated information is, in turn, a subset of

$$(Y^\delta(T_{\text{left}}) - (K_1 + 1)\delta, Y^\delta(T_{\text{left}}) + (K_1 + 1)\delta).$$

Here, recall that $K_1 := (1 - \alpha)^{-1}$ is the Lipschitz constant of the Skorokhod reflection map $S.$ Consecutively, if we propose a sample Z from the uniform distribution in the interval $(Y^\delta(T_{\text{left}}) - (K_1 + 1)\delta, Y^\delta(T_{\text{left}}) + (K_1 + 1)\delta),$ then the likelihood ratio (or the Radon–Nikodym derivative) between the true conditional density and proposal density is proportional to $f_\Delta(w - Y(T_{\text{left}}));$ consecutively, a traditional accept/reject algorithm would accept the proposed sample Z if

$$V < L(Z; Y(T_{\text{left}})) := \frac{\pi(Z - Y(T_{\text{left}}))\rho(Z - Y(T_{\text{left}}))}{c_\pi c_\rho}$$

for an independent $V \sim \text{uniform}(0, 1);$ hereafter, we use $L(z; y)$ to denote

$$L(z; y) := \frac{\pi(z - y)\rho(z - y)}{c_\rho c_\phi}.$$

However, as we do not know $Y(T_{\text{left}})$ exactly, if $L(Z; \cdot)$ is Lipschitz continuous, the fact that $Y(T_{\text{left}})$ differs from $Y^\delta(T_{\text{left}})$ by $\pm K_1\delta$ can be used to implement a refine-until-accept/reject sampler discussed as in Section 2.1.

The Lipschitz continuity of $L(z; y),$ as a function of $y,$ follows as a simple consequence of Lemma 3.1 established earlier. If f and g are Lipschitz continuous with Lipschitz constants K_f and $K_g,$ and respective absolute bounds c_f and $c_g,$ then fg is a Lipschitz continuous function with Lipschitz constant at most $c_f K_g + c_g K_f.$ As a result, the product $\pi(\cdot) \times \rho(\cdot)$ is Lipschitz

continuous as well with Lipschitz constant $c_\pi K_\rho + c_\rho K_\pi$. Consequently, the function $L(z, y)$ is Lipschitz continuous, as a function of y , with Lipschitz constant

$$K_2 := \frac{K_\pi}{c_\pi} + \frac{K_\rho}{c_\rho}.$$

Given this Lipschitz continuity of L , one can unambiguously accept the proposal Z if

$$V < L(Z; Y^\delta(T_{\text{left}})) - K_1 K_2 \delta,$$

or reject the proposal Z conclusively if

$$V > L(Z; Y^\delta(T_{\text{left}})) + K_1 K_2 \delta.$$

However, if V is within $L(Z; Y^\delta(T_{\text{left}})) \pm K_1 K_2 \delta$, then we obtain a more refined approximation $Y^{\varepsilon_n}(T_{\text{left}})$, for a suitable $\varepsilon_n < \delta$, that is good enough to decide whether to accept/reject. In particular, if ε_n is smaller than $(K_1 K_2)^{-1} |V - L(Z; Y^{\varepsilon_n}(T_{\text{left}}))|$, then

$$V < L(Z; Y(T_{\text{left}})) \iff V < L(Z; Y^{\varepsilon_n}(T_{\text{left}})) - K_1 K_2 \varepsilon_n. \tag{3.6}$$

This equivalent comparison is at the heart of the refine-until-accept/reject sampler in Algorithm 3 below. It takes the intersection layers \mathcal{I} returned by Algorithm 2 as input, and generates further refined approximations $(Y^{\varepsilon_n}(t) : t \in [0, T_{\text{left}}])$ of the RBM, if necessary, in order to perform the equivalent comparisons in (3.6).

Algorithm 3. To accept/reject the proposal $Z \sim \text{uniform}(Y^\delta(T_{\text{left}}) - (K_1 + 1)\delta, Y^\delta(T_{\text{left}}) + (K_1 + 1)\delta)$. If $V < L(Z; Y(T_{\text{left}}))$, the algorithm returns Z ; otherwise it rejects Z and returns nothing. Recall that $L(z; y) := c_\pi^{-1} c_\rho^{-1} \pi(z - y) \rho(z - y)$, $K_1 := 1/(1 - \alpha)$, and $K_2 := K_\pi/c_\pi + K_\rho/c_\rho$.

procedure RARsampler($\mathcal{I}, T_{\text{left}}, N, \delta, Y^\delta(T_{\text{left}}), Z$)

Initialize $n = N, \varepsilon_n = \delta$

Draw V uniformly from $[0, 1]$

while $\varepsilon_n > (K_1 K_2)^{-1} |V - L(Z; Y^{\varepsilon_n}(T_{\text{left}}))|$ **do**

Increment $n \leftarrow n + 1$

Simulate the intersection layers $\mathcal{I}_{j,n}$ for $j = 1, \dots, 2^n$ conditional on \mathcal{I}

Form $B_n(\cdot)$ as in (3.1), which serves as piecewise linear approximation to $B(\cdot)$

Set $\varepsilon_n = \max\{U_{j,n}^\uparrow - L_{j,n}^\downarrow : j = 1, \dots, 2^n\}$ and $\mathcal{I} = \{\mathcal{I}_{j,n} : j = 1, \dots, 2^n\}$

Letting $X(t) = B_n(t)$ for $t \in [0, T_{\text{left}}]$ in (1.1), solve for the reflected process $Y(\cdot)$; call the solution as $Y_n(\cdot)$; set the required refined approximation $Y^{\varepsilon_n}(T_{\text{left}}) = Y_n(T_{\text{left}})$.

end while

if $V < L(Z; Y^{\varepsilon_n}(T_{\text{left}})) - K_1 K_2 \varepsilon_n$ **then**

Return Z

else

Return \emptyset

end if

For d -dimensional processes, the probability density of the increment Δ and proposal density $g(\cdot)$ are both given by the product of one-dimensional densities. This results in a likelihood ratio which is also of product form, leading to a straightforward generalization of the refine-until-accept/reject procedure given in Algorithm 3.

4. A note on computational complexity

Our objective in this section is to understand the computational effort required to execute the refine-until-accept/reject exact sampler described in Section 3. For ease of exposition, we do not keep track of multiplying constants but instead adopt the following standard notation to describe the asymptotic behavior of functions. For given functions $f: \mathbb{R}^+ \rightarrow \mathbb{R}^+$ and $g: \mathbb{R}^+ \rightarrow \mathbb{R}^+$, we say $f(x) = O(g(x))$ if there exists $c_1 > 0$ and large enough x_1 such that $f(x) \leq c_1 g(x)$ for all $x > x_1$; further, we say $f(x) = \Theta(g(x))$ if there also exists $c_2 > 0$ and large enough x_2 such that $c_2 g(x) \leq f(x) \leq c_1 g(x)$ for all $x > x_2$.

Recall the definition of the error in approximation ε_n in (3.2), and to achieve this accuracy we needed to simulate relevant information (such as maxima, minima, and endpoints for d independent one-dimensional Brownian motions $B_i(\cdot)$ in 2^n dyadic intervals as described in Section 3.1.1. As this entails a constant amount of expected computation for each dimension in each of the 2^n dyadic intervals, the computational cost at the end of n steps of the iterative procedure described in Section 3.1.1 is $\Theta(d2^n)$.

Next, observe that the preconditioning step requires us to iterate until the stopping time $N = \inf\{n \geq 1: Y_{n,i}(t) > (K_1 + 1)\varepsilon_n, i = 1, \dots, d\}$. As the computational cost of solving the Skorokhod problem (as in Algorithm 1) with a piecewise linear input $B_k(\cdot)$ is $\Theta(d^3 2^k)$ uniformly in d , the total cost of executing the entire preconditioning step is $\Theta(d^3 \sum_{k=1}^N 2^k) = \Theta(d^3 2^N)$. Here, we have used the fact that $O(d^3)$ computations are needed to solve the linear program in Algorithm 1 when the set C contains $O(d)$ elements. Following the same line of reasoning, if we let

$$N' = \inf\{n \geq N: \varepsilon_n < (K_1 K_2)^{-1} |V - L(\mathbf{Z}; \mathbf{Y}^{\varepsilon_n}(T_{\text{left}}))|\},$$

as required in the refine-until-accept/reject step in Algorithm 3, the corresponding computational cost is $\Theta(d^3 2^{N'})$. As $N' \geq N$, the total computational cost of the sampling procedure is $\Theta(d^3 2^N + d^3 2^{N'}) = \Theta(d^3 2^{N'})$. In other words, there exist positive constants c_1 and c_2 such that

$$c_1 d^3 2^{N'} \leq \text{‘computational cost’} \leq c_2 d^3 2^{N'}.$$

Therefore, the expected computational cost of the entire sampling procedure is $\Theta(d^3 \mathbb{E}[2^{N'}])$.

Next, to compute $\mathbb{E}[2^{N'}]$, we first use the definition of ε_n in (3.2) to observe that

$$\mathbb{P}(N' > n) = \mathbb{P}\left(\varepsilon_n > \frac{D}{K_1 K_2}\right) = \mathbb{P}\left(\max_{i=1, \dots, d} \sup_{t \in [0, 1]} |B_{n,i}^\uparrow(t) - B_{n,i}^\downarrow(t)| > \frac{D}{K_1 K_2}\right),$$

where we let $D = |V - L(\mathbf{Z}; \mathbf{Y}^{\varepsilon_n}(T_{\text{left}}))|$ for notational convenience. For each fixed i , it follows from the construction of one-dimensional piecewise constant bounding processes $B_{n,i}^\uparrow(\cdot)$ and

$B_{n,i}^\downarrow$ in Section 3.1.1 that

$$\begin{aligned} \sup_{t \in [0,1]} |B_{n,i}^\uparrow(t) - B_{n,i}^\downarrow(t)| &= \max_{j=1, \dots, 2^n} |U_{j,n}^\uparrow - L_{j,n}^\downarrow| \\ &\leq \max_{j=1, \dots, 2^n} \{|U_{j,n}^\uparrow - M_{j,n}| + |M_{j,n} - m_{j,n}| + |m_{j,n} - L_{j,n}^\downarrow|\} \\ &\leq \max_{j=1, \dots, 2^n} \{2^{-(n+1)/2} + |M_{j,n} - m_{j,n}| + 2^{-(n+1)/2}\} \\ &\stackrel{D}{=} 2^{-n/2} + \max_{j=1, \dots, 2^n} \left(\sup_{t \in [0,2^n]} W_j(t) - \inf_{t \in [0,2^n]} W_j(t) \right), \end{aligned}$$

where $(W_j(t) : j = 1, \dots, 2^n)$ are 2^n independent copies of standard Brownian motion. Here, the notation ‘ $\stackrel{D}{=}$ ’ is used to denote equality in distribution. If we let

$$\bar{Z}_j := 1 + \sup_{t \in [0,1]} W_j(t) - \inf_{t \in [0,1]} W_j(t),$$

then due to self-similarity of Brownian motion,

$$\sup_{t \in [0,1]} |B_{n,i}^\uparrow(t) - B_{n,i}^\downarrow(t)| \stackrel{D}{=} 2^{-n/2} + 2^{-n/2} \left(\sup_{t \in [0,1]} W_j(t) - \inf_{t \in [0,1]} W_j(t) \right) = 2^{-n/2} \max_{j=1, \dots, 2^n} \bar{Z}_j.$$

Since the approximations $B_{n,i}$ are independently obtained for each $i = 1, \dots, d$, we have

$$\sup_{t \in [0,1]} |B_{n,i}^\uparrow(t) - B_{n,i}^\downarrow(t)| = \max_{k=1, \dots, d2^n} \bar{Z}_k,$$

where $(\bar{Z}_k : k = 1, \dots, d2^n)$ are $d2^n$ independent copies of $\bar{Z} := 1 + \sup_{t \in [0,1]} W(t) - \inf_{t \in [0,1]} W(t)$. For notational convenience, denote $M_n = \max_{j=1, \dots, d2^n} \bar{Z}_j$. Then

$$\begin{aligned} \mathbb{P}(N' > n) &= \mathbb{P}\left(2^{-n/2} M_n > \frac{D}{K_1 K_2}\right) \\ &= \mathbb{P}\left(D < K_1 K_2 \frac{M_n}{2^{n/2}}\right) \\ &= \mathbb{P}\left(L(\mathbf{Z}; \mathbf{Y}^{\varepsilon_n}(T_{\text{left}})) - K_1 K_2 \frac{M_n}{2^{n/2}} < V < L(\mathbf{Z}; \mathbf{Y}^{\varepsilon_n}(T_{\text{left}})) + K_1 K_2 \frac{M_n}{2^{n/2}}\right) \end{aligned}$$

due to our earlier definition that $D = |V - L(\mathbf{Z}; \mathbf{Y}^{\varepsilon_n}(T_{\text{left}}))|$. As $V \sim \text{uniform}[0, 1]$, it is immediate that

$$\mathbb{P}(N' > n) = \frac{2K_1 K_2}{2^{n/2}} \mathbb{E}[M_n] = \Theta\left(\sqrt{\frac{n}{2^n}}\right),$$

where the second equality follows from the observation that $\mathbb{E}[M_n] = \Theta(\sqrt{n})$, which is proved in Lemma A.2. Therefore, $\mathbb{P}(2^{N'} > x) = \Theta(x^{-0.5} \sqrt{\log x})$. As the random variable $2^{N'}$ has regularly varying tails with index -0.5 ,

$$\mathbb{E}[2^{N'}] = \infty,$$

and, consequently, the expected total computational cost is infinite.

An alternative, intuitive explanation for why the expected termination time is infinite is as follows. Note that conditional on \mathbf{Z} and $\mathbf{Y}^{\varepsilon_n}(T_{\text{left}})$, the distance $D = |V - L(\mathbf{Z}; \mathbf{Y}^{\varepsilon_n}(T_{\text{left}}))|$

is less than δ with probability $O(\delta)$ (because V is uniformly distributed). Thus, if the cost of generating Y^D (required for the decision whether to accept or reject) is $C(D)$, the running time of the algorithm would be finite if $\int_0^1 C(u) du < \infty$. Unfortunately, however, the cost of producing an ε -strong approximation to Brownian motion (X^ε) is roughly $O(1/\varepsilon^2)$ (see, for example, [4]) and, therefore, $C(D) \geq c/D^2$ with positive probability for some $c > 0$, which results in an infinite expected running time.

5. Conclusions

We provide the first exact sampling algorithm to obtain samples from a multidimensional reflected Brownian motion. The algorithm relies on a novel conditional acceptance–rejection step, which is implemented by carefully refining ε -strong approximations of the reflected Brownian motion path $Y(\cdot)$ until we can conclusively accept or reject a proposal Z from a suitable uniform distribution. Unfortunately, as shown in Section 4, the proposed algorithm has infinite expected termination time due to the large amount of computational effort required to conclusively decide whether $V < L(Z; Y(T_{\text{left}}))$ when the proposal likelihood $L(Z; Y(T_{\text{left}}))$ and the uniform random variable V are close. It may be of interest to the reader to know whether the entire exact sampling scheme can be executed with finite expected computational effort if we, somehow, are able to resolve the difficulty in deciding whether $V < L(Z; Y(T_{\text{left}}))$ with finite expected computational effort. We believe this is indeed the case because of the following reasoning.

Apart from the refine-until-accept/reject step in Algorithm 3, the only other step where we execute a ‘while’ loop performing a random comparison is in the preconditioning procedure in Algorithm 2. Recall that the preconditioning step must develop a piecewise approximation to Brownian motion that is accurate enough to satisfy $\varepsilon_n := \|B_n(\cdot) - B(\cdot)\| < (K_1 + 1)Y_i(T)$ for $i = 1, \dots, d$, in order to identify $(T_{\text{left}}, T_{\text{right}})$. If the probability density of $Y_i(T)$ evaluated at 0 is positive (as in the one-dimensional RBM case), then the probability that $Y_i(T)$ is smaller than δ is at least $c_1\delta$ (for every $\delta < \delta'$ suitably small), and the computational effort required to generate a Brownian approximation that satisfies $\varepsilon_n < (K_1 + 1)\delta$ is larger than c_2/δ^2 with positive probability; here, c_1 and c_2 are suitable positive constants. As the required computational effort $O(1/(\min_i Y_i(T))^2)$ is high when $Y_i(T)$ is close to 0 for some i , the expected computational effort required in the preconditioning step is at least

$$\int_{c_2/\delta^2}^\infty \mathbb{P}(\text{computational cost} > u) du \geq \int_{c_2/\delta^2}^\infty c_1 \sqrt{\frac{c_2}{u}} du = \infty.$$

However, this difficulty can be easily resolved if we imagine, for a moment, that it is possible to resolve the earlier difficulty explained in Section 4 (on deciding whether $V < L(Z; Y(T_{\text{left}}))$ within finite expected time), and it is possible to obtain exact samples of $Y(T)$ whenever $Y(T)$ is bounded away from the reflecting barrier; specifically, let us assume we can obtain samples of $Y(T)$ with finite expected computational effort C_γ when $\min_i Y_i(T) > \gamma$ for some fixed constant $\gamma \in (0, 1)$. In that case, we first obtain an exact sample of $Y(t')$ for the latest $t' \leq T$ during which $\min_i Y_i(t') > \gamma$. To be specific, define $t' := \sup\{t \leq T : \min_i Y_i(t) > \gamma\}$, $\tilde{Y} := Y_i(t')$, and $T_\gamma := T - t'$. Then, due to the Markov property and self-similarity of RBM, the original objective of obtaining a sample of $Y(T)$ can be equivalently written as follows. Obtain a sample of $Y_{\text{new}}(T_\gamma/\gamma^2)$, where $Y_{\text{new}}(\cdot)$ is also an RBM obtained by shifting and scaling the RBM $Y(\cdot)$ as in $Y_{\text{new}}(t) := Y(t' + \gamma^2 t)/\gamma$ with initial condition $Y_{\text{new}}(0) := \tilde{Y}/\gamma$.

Thus, even if the original problem of detecting $(T_{\text{left}}, T_{\text{right}})$ is difficult when $Y(T) < \gamma$, by a suitable translation and scaling (magnification) of the underlying Brownian and RBM paths, we have a new, but equivalent, objective of sampling from $Y_{\text{new}}(T_\gamma/\gamma^2) = Y(T)/\gamma$. In the case when $Y_{\text{new}}(T_\gamma/\gamma^2)$ is smaller than γ as well, we perform a similar translation and scaling once again recursively. Since $\mathbb{E}[T_\gamma/\gamma^2]$ is uniformly bounded as a function of γ , a simple recursive algorithm complexity analysis (see, for example, [8, Chapter 4]) yields that the total expected computational effort of the described recursive procedure is $O(C_\gamma/\mathbb{P}(\min_i Y_i(T) > \gamma))$, which is finite, as per our assumption on C_γ .

As the described recursive construction for the preconditioning step is built on the assumption that we can conclusively decide whether $V < L(\mathbf{Z}; Y(T_{\text{left}}))$ within finite expected computational effort, we identify the difficulty explained carefully in Section 4 as the only fundamental bottleneck in obtaining exact samples of multidimensional RBM. Future research that addresses this bottleneck by means of new techniques will make the proposed algorithm, which is currently of theoretical importance, to be more suitable for practice as well.

Appendix A

Here we provide the proof of Lemma 3.1, and present explicit expressions for the constants $K_\pi, K_\rho, c_\pi,$ and c_ρ . In order to prove Lemma 3.1, we need the following result.

Lemma A.1. *For any given $U > L, r > 0$, the function $\gamma(L, U; r, a, b)$ defined in (3.5) is Lipschitz continuous with respect to the variables a and b ; that is,*

$$|\gamma(L, U; r, a_1, \cdot) - \gamma(L, U; r, a_2, \cdot)| < K(L, U, r)|a_1 - a_2|,$$

$$|\gamma(L, U; r, \cdot, b_1) - \gamma(L, U; r, \cdot, b_2)| < K(L, U, r)|b_1 - b_2|$$

for all $a_1, a_2, b_1, b_2 \in (L, U)$. The Lipschitz constant $K(L, U, r)$ is given by

$$K(L, U, r) := \sum_{j \geq 1} K_j = \frac{8(U - L)}{r} \sum_{j \geq 1} j \exp\left(-\frac{2}{r}(U - L)^2(j - 1)^2\right).$$

Proof. Let $\gamma_n(a, b) = 1 - \sum_{j=1}^n (\sigma_j - \tau_j)$. Since a, b take values in (L, U) , it is easily checked that for all $j \geq 1$ both

$$\left| \frac{d}{da}(\sigma_j - \tau_j) \right| < K_j \quad \text{and} \quad \left| \frac{d}{db}(\sigma_j - \tau_j) \right| < K_j,$$

where

$$K_j := \frac{8(U - L)j}{r} \exp\left(-\frac{2}{r}(U - L)^2(j - 1)^2\right).$$

Then it is immediate that, for all n ,

$$\left| \frac{d\gamma_n(a, b)}{da} \right| < \sum_{j=1}^\infty K_j \quad \text{and} \quad \left| \frac{d\gamma_n(a, b)}{db} \right| < \sum_{j=1}^\infty K_j.$$

As a consequence, we use the following elementary properties of Lipschitz continuity to establish the Lipschitz continuity of $\gamma(L, U; r, a, b)$ with respect to variables a and b .

- If a differentiable function $f(\cdot)$ on a convex domain is such that its first derivative $|f'(x)| < K$ for some constant K , then the function $f(\cdot)$ is Lipschitz continuous with Lipschitz constant at most K .

- If a sequence of Lipschitz functions $f_n(\cdot)$ all having Lipschitz constant bounded by K converge uniformly to $f(\cdot)$, then $f(\cdot)$ is also Lipschitz continuous with Lipschitz constant at most K .

Since $\gamma_n(a, b)$ converge uniformly to $\gamma(a, b)$ for $a, b \in (L, U)$, it follows immediately from the above two facts that $\gamma(L, U; r, a, b)$ is Lipschitz continuous with Lipschitz constant at most

$$K(L, U, r) := \sum_{j \geq 1} K_j = \frac{8(U - L)}{r} \sum_{j \geq 1} j \exp\left(-\frac{2}{r}(U - L)^2(j - 1)^2\right). \quad \square$$

Proof of Lemma 3.1. For all $x \in (L^\downarrow, U^\uparrow)$,

$$\pi(x) = \exp\left(-\frac{1}{2}\left(x - \frac{s}{l}v\right)^2 \Big/ \frac{s(l - s)}{l}\right) \leq 1 =: c_\pi. \quad (\text{A.1})$$

The Lipschitz continuity of $\pi(\cdot)$ follows from the boundedness of its first derivative $\pi'(\cdot)$ on the convex domain (L, U) . For all $x \in (L, U)$,

$$\left| \frac{d\pi(x)}{dx} \right| \leq \frac{|xl - sv|}{s(l - s)} < \frac{\max\{|Ul - sv|, |Ll - sv|\}}{s(l - s)} =: K_\pi. \quad (\text{A.2})$$

To prove the Lipschitz continuity of $\rho(\cdot)$, we first note the boundedness of $\gamma(L, U; r, \cdot, \cdot)$. Simple substitution will yield that $\gamma(L, U; r, a, b) = 0$ whenever either a or b equals one of L or U . Then due to the Lipschitz continuity of $\gamma(L, U; r, \cdot, \cdot)$, from Lemma A.1, we have

$$|\gamma(L, U; r, a, b)| \leq K(L, U, r)(U - L). \quad (\text{A.3})$$

Now consider the first term $\gamma_1(x)\gamma_2(x)$ in (3.4).

- Due to (A.3), $|\gamma_1(\cdot)|$ and $|\gamma_2(\cdot)|$ are bounded by $K(L^\downarrow, U^\uparrow, s)(U^\uparrow - L^\downarrow)$ and $K(L^\downarrow, U^\uparrow, l - s)(U^\uparrow - L^\downarrow)$, respectively, in the interval $x \in (L^\downarrow, U^\uparrow)$.
- From Lemma A.1, it follows that $\gamma_1(\cdot)$ and $\gamma_2(\cdot)$ are Lipschitz continuous (with respect to the variable x) with Lipschitz constants at most $K(L^\downarrow, U^\uparrow, s)$ and $K(L^\downarrow, U^\uparrow, l - s)$, respectively.

From the above two observations, we conclude that $\gamma_1(\cdot)\gamma_2(\cdot)$ is Lipschitz continuous with respect to x with Lipschitz constant at most

$$K_{1,2} := 2K(L^\downarrow, U^\uparrow, s)K(L^\downarrow, U^\uparrow, l - s)(U^\uparrow - L^\downarrow).$$

This is because if f, g are Lipschitz continuous with respective Lipschitz constants K_f and K_g and absolute bounds C_f and C_g , then fg is Lipschitz continuous with Lipschitz constant at most $C_f K_g + C_g K_f$. Using the same reasoning, the Lipschitz constants of other terms in (3.4), namely, $\gamma_3(\cdot)\gamma_4(\cdot)$, $\gamma_5(\cdot)\gamma_6(\cdot)$, and $\gamma_7(\cdot)\gamma_8(\cdot)$ are at most

$$\begin{aligned} K_{3,4} &:= 2K(L^\uparrow, U^\uparrow, s)K(L^\uparrow, U^\uparrow, l - s)(U^\uparrow - L^\uparrow), \\ K_{5,6} &:= 2K(L^\downarrow, U^\downarrow, s)K(L^\downarrow, U^\downarrow, l - s)(U^\downarrow - L^\downarrow), \end{aligned}$$

and less than

$$K_{7,8} := 2K(L^\uparrow, U^\downarrow, s)K(L^\uparrow, U^\downarrow, l - s)(U^\downarrow - L^\uparrow),$$

respectively. Therefore, $\rho(x)$ is Lipschitz continuous with Lipschitz constant K_ρ given by

$$K_\rho := K_{1,2} + K_{3,4} + K_{5,6} + K_{7,8}.$$

Since $\rho(x) = 0$ whenever x takes either L^\downarrow or U^\uparrow , using Lipschitz continuity of ρ , we reason that

$$|\rho(x)| \leq K_\rho(U^\uparrow - L^\downarrow) =: c_\rho.$$

This along with (A.1) and (A.2) proves the claim. □

Lemma A.2. *Recall the definition $\bar{Z} := 1 + \sup_{t \in [0,1]} W(t) - \inf_{t \in [0,1]} W(t)$, where $(W(t) : t \in [0, 1])$ is a standard Brownian motion. If $(Z_i : i = 1, \dots, k)$ are k independent copies of \bar{Z} , then*

$$\mathbb{E} \left[\max_{i=1, \dots, k} \bar{Z}_i \right] = \Theta(\sqrt{\log k}) \quad \text{as } k \rightarrow \infty.$$

Proof. We first observe that

$$\begin{aligned} \max_{i=1, \dots, k} \bar{Z}_i &= \max_{i=1, \dots, k} \left(1 + \sup_{t \in [0,1]} W(t) - \inf_{t \in [0,1]} W(t) \right) \\ &\leq 1 + \max_{i=1, \dots, k} \sup_{t \in [0,1]} W(t) + \max_{i=1, \dots, k} \left(- \inf_{t \in [0,1]} W(t) \right). \end{aligned}$$

As $\sup_{t \in [0,1]} W(t) \stackrel{D}{=} - \inf_{t \in [0,1]} W(t)$, we have

$$1 + \max_{i=1, \dots, k} \sup_{t \in [0,1]} W(t) \stackrel{D}{\leq} \max_{i=1, \dots, k} \bar{Z}_i \stackrel{D}{\leq} 1 + 2 \max_{i=1, \dots, k} \sup_{t \in [0,1]} W(t),$$

where the notation ' $X \stackrel{D}{\leq} Y$ ' denotes that X is stochastically upper bounded by Y . Further, as $\sup_{t \in [0,1]} W(t) \stackrel{D}{=} |Z|$ when Z follows a standard normal distribution, we have

$$\mathbb{E} \left[\max_{i=1, \dots, k} \bar{Z}_i \right] = \Theta \left(\mathbb{E} \left[\max_{i=1, \dots, k} \sup_{t \in [0,1]} W(t) \right] \right) = \Theta \left(\mathbb{E} \left[\max_{i=1, \dots, k} |Z_i| \right] \right) \quad \text{as } k \rightarrow \infty. \quad (\text{A.4})$$

Here, $(Z_i : i = 1, \dots, k)$ are simply k independent copies of a standard normal variable. Next, if we denote the positive and negative parts of Z_i as $Z_i^+ = \max\{Z_i, 0\}$ and $Z_i^- = -\min\{Z_i, 0\}$, then $|Z_i| = Z_i^+ + Z_i^-$. Further, as

$$\max_{i=1, \dots, k} Z_i^+ \leq \max_{i=1, \dots, k} |Z_i| \leq \max_{i=1, \dots, k} Z_i^+ + \max_{i=1, \dots, k} Z_i^-,$$

it follows from (A.4) that

$$\mathbb{E} \left[\max_{i=1, \dots, k} \bar{Z}_i \right] = \Theta \left(\mathbb{E} \left[\max_{i=1, \dots, k} Z_i^+ \right] \right) = \Theta(\mathbb{E}[M^+]),$$

where M^+ is the positive part of $M := \max_{i=1, \dots, k} Z_i$. Since $\mathbb{E}[M] = \Theta(\sqrt{\log k})$ and $\mathbb{E}[M^-] = \int_0^\infty (\mathbb{P}(Z < -u))^k du \rightarrow 0$ as $k \rightarrow \infty$, we have $\mathbb{E}[M^+] = \mathbb{E}[M] + \mathbb{E}[M^-] = \Theta(\sqrt{\log k})$, thus proving the claim. □

Acknowledgement

J. Blanchet gratefully acknowledges the support of the National Science Foundation (grant number 1538217).

References

- [1] ASMUSSEN, S. AND GLYNN, P. W. (2007). *Stochastic Simulation: Algorithms and Analysis*. Springer, New York.
- [2] BESKOS, A. AND ROBERTS, G. O. (2005). Exact simulation of diffusions. *Ann. Appl. Prob.* **15**, 2422–2444.
- [3] BESKOS, A., PAPASPILIOPOULOS, O. AND ROBERTS, G. O. (2006). Retrospective exact simulation of diffusion sample paths with applications. *Bernoulli* **12**, 1077–1098.
- [4] BESKOS, A., PELUCHETTI, S. AND ROBERTS, G. (2012). ε -strong simulation of the Brownian path. *Bernoulli* **18**, 1223–1248.
- [5] BLANCHET, J. AND CHEN, X. (2015). Steady-state simulation of reflected Brownian motion and related stochastic networks. *Ann. Appl. Prob.* **25**, 3209–3250.
- [6] CHEN, H. AND YAO, D. D. (2001). *Fundamentals of Queueing Networks: Performance, Asymptotics, and Optimization* (Appl. Math. (New York) **46**). Springer, New York.
- [7] CHEN, N. AND HUANG, Z. (2013). Localization and exact simulation of Brownian motion-driven stochastic differential equations. *Math. Operat. Res.* **38**, 591–616.
- [8] CORMEN, T. H., LEISERSON, C. E., RIVEST, R. L. AND STEIN, C. (2001). *Introduction to Algorithms*, 2nd edn. McGraw-Hill, Boston, MA.
- [9] ÉTORÉ, P. AND MARTINEZ, M. (2013). Exact simulation of one-dimensional stochastic differential equations involving the local time at zero of the unknown process. *Monte Carlo Methods Appl.* **19**, 41–71.
- [10] HARRISON, J. M. AND REIMAN, M. I. (1981). Reflected Brownian motion on an orthant. *Ann. Prob.* **9**, 302–308.
- [11] KELLA, O. AND WHITT, W. (1996). Stability and structural properties of stochastic storage networks. *J. Appl. Prob.* **33**, 1169–1180.
- [12] POLLOCK, M., JOHANSEN, A. M. AND ROBERTS, G. O. (2016). On the exact and ε -strong simulation of (jump) diffusions. *Bernoulli* **22**, 794–856.
- [13] PÖTZELBERGER, K. AND WANG, L. (2001). Boundary crossing probability for Brownian motion. *J. Appl. Prob.* **38**, 152–164.
- [14] REIMAN, M. I. (1984). Open queueing networks in heavy traffic. *Math. Operat. Res.* **9**, 441–458.